

THE ASTEROSEISMOLOGY METACOMPUTER

T.S. Metcalfe and R.E. Nather

Department of Astronomy, University of Texas, Austin, TX 78701 U.S.A.

Received September 1, 1999.

Abstract. We have developed a specialized computational instrument for fitting models of pulsating white dwarfs to observations made with the Whole Earth Telescope. This metacomputer makes use of inexpensive PC hardware and free software, including a parallel genetic algorithm which performs a global search for the best-fit set of parameters.

Key words: instrumentation: miscellaneous – methods: numerical – stars: white dwarfs

1. INTRODUCTION

White dwarf asteroseismology offers the opportunity to probe the structure and composition of stellar objects governed by relatively simple principles. The observational requirements of asteroseismology have been addressed by the development of the Whole Earth Telescope (WET), but the analytical procedures need to be refined to take full advantage of the possibilities afforded by the WET data.

The adjustable parameters in our computer models of white dwarfs presently include the total mass, the temperature, the H and He layer masses, the core composition, and the transition zone thicknesses. Finding a proper set of these to provide a close fit to the observed data is difficult. The current procedure is a cut-and-try process guided by intuition and experience, and is far more subjective than we would like. Objective procedures for determining the best-fit model are essential if asteroseismology is to become a widely-accepted and reliable astronomical technique. We must be able to demonstrate that, within the range of different values the model parameters can assume, we have found the only solution, or

the best one if more than one is possible. To address this problem, we are applying a search-and-fit technique employing a genetic algorithm (GA), which can explore the myriad parameter combinations possible and select for us the best one, or ones (cf. Goldberg 1989, Charbonneau 1995, Metcalfe & Nather 1999).

Although genetic algorithms are more efficient than other comparably global techniques, they are still quite demanding computationally. To be practical, the GA-based fitting technique requires a dedicated instrument to perform the calculations. Over the past year, we have designed and configured such an instrument—an isolated network of 64 minimal PCs running Linux. Since the structure of a GA is very conducive to parallelization, this metacomputer allows us to run our code much faster than would otherwise be possible.

2. HARDWARE

In January 1998, around the time that the idea of commodity parallel processing started getting a lot of attention, we were independently designing a metacomputer of our own. Our budget was modest, so we set out to get the best performance possible per dollar without restricting the ability of the machine to solve our specific problem.

The original Beowulf cluster (Becker et al. 1995), which we didn't know about at the time, had a number of features which, though they contributed to the utility of the machine as a multi-purpose computational tool, were unnecessary for our particular problem. We wanted to use each node of the metacomputer to run identical tasks with small, independent sets of data. The results of the calculations performed by the nodes consisted of just a few numbers which only needed to be communicated to the master process, never to another node. Essentially, network bandwidth was not an issue because the computation to communication ratio of our application was extremely high, and hard disks were not needed on the nodes because our problem did not require any significant amount of data swapping. In the end we settled on a design including one master server augmented by minimal nodes connected by a simple 10base-2 network (see Figure 1).

The master computer is a Pentium-II 333 MHz system with three NE-2000 compatible network cards, each of which drives 1/3 of the nodes on a subnet. Since a single ethernet card can handle up to 30 devices, no repeater was necessary.



Fig. 1. The 64 minimal nodes of the metacomputer on shelves surrounding the master computer.

The slave nodes were assembled from components obtained at a discount computer outlet. Each node consists of an ATX tower case with a motherboard, processor and fan, a single 32 MB SDRAM, and an NE-2000 compatible network card with a custom made boot-EPROM. The nodes are connected in series with 3-ft ethernet coaxial cables. Half of the nodes contain Pentium-II 300 MHz processors, while the other half are AMD K6-II 450 MHz chips. The total cost of the system was around \$25k, but it could be built for considerably less today, and less still tomorrow.

3. SOFTWARE

To make the metacomputer work, we relied on the open-source Linux operating system and software tools. We programmed the EPROMs with Gero Kuhlmann's NETBOOT package to allow each node to download and mount an independent Linux filesystem on a small ramdisk partition. We used Tom Fawcett's YARD package to create the filesystem, and we included in it a pared down version of the PVM software developed at Oak Ridge National Laboratory (Geist et al. 1994).

Finally, we incorporated the message passing routines of the PVM library into PIKAIA, a general purpose public-domain GA developed by Charbonneau (1995), and we modified the white dwarf evolution and pulsation codes (see Wood 1990, Bradley 1993, Mont-

gomery 1998) to allow reliable and automated calculation of the normal modes of oscillation for white dwarf stars with a wide range of masses, temperatures, and other parameters.

4. BENCHMARKS

Measuring the absolute performance of the metacomputer is difficult because the result strongly depends on the fraction of Floating-point Division operations (FDIVs) used in the benchmark code. Table 1 lists four different measures of the absolute speed in Millions of Floating-point Operations Per Second (MFLOPS).

Table 1. The absolute speed of the metacomputer.

| Benchmark | P-II 300 MHz | K6-II 450 MHz | Total Speed |
|-----------|--------------|---------------|-------------|
| MFLOPS(1) | 80.6 | 65.1 | 4662.4 |
| MFLOPS(2) | 47.9 | 67.7 | 3699.2 |
| MFLOPS(3) | 56.8 | 106.9 | 7056.0 |
| MFLOPS(4) | 65.5 | 158.9 | 7180.8 |

The code for MFLOPS(1) is essentially scalar—that is, vector processor performance will reflect scalar performance which will lie far below expected vector performance. Also, the percentage of FDIVs (9.6%) is considered somewhat high. The code for MFLOPS(2) is fully vectorizable. The percentage of FDIVs (9.2%) is still somewhat on the high side. The code for MFLOPS(3) is also fully vectorizable. The percentage of FDIVs (3.4%) is considered moderate. The code for MFLOPS(4) is fully vectorizable, but the percentage of FDIVs is zero.

We feel that MFLOPS(3) provides the best measure of the expected performance for the white dwarf code, because of the moderate percentage of FDIVs. Adopting this value, we have achieved a price to performance ratio near \$3.50/MFLOPS.

ACKNOWLEDGEMENTS. We would like to thank Gary Hansen for donating the 32 K6-II 450 MHz processors through an arrangement with AMD. This work was made possible by a grant from the National Science Foundation.

REFERENCES

- Becker, D., Sterling, T., Savarese, D., Dorband, J., Ranawak, U., and Packer, C. 1995, Beowulf: A Parallel Workstation for Scientific Computation, in Proceedings of the International Conference on Parallel Processing (New York: Institute of Electrical and Electronics Engineers).
- Bradley, P. 1993, Ph.D. Thesis, University of Texas at Austin.
- Charbonneau, P. 1995, ApJS, 101, 309.
- Geist, A., Beguelin, A., Dongarra, J., Jiang, W., Manchek, R., and Sunderam, V. 1994, PVM: Parallel Virtual Machine, A Users' Guide and Tutorial for Networked Parallel Computing, (Cambridge: MIT Press).
- Goldberg, D. 1989, Genetic Algorithms in Search, Optimization, and Machine Learning, (Reading, MA: Addison Wesley).
- Metcalfe, T.S. and Nather, R.E. 1999, Linux Journal, 65, 58.
- Montgomery, M. 1998, Ph.D. Thesis, University of Texas at Austin.
- Wood, M. 1990, Ph.D. Thesis, University of Texas at Austin.